Supporting Incremental Knowledge Elicitation in Decision-Theoretic Systems

Proceedings of the AAAI Spring Symposium Workshop on Interactive and Mixed-Initiative Decision-Theoretic Systems, pp. 14-15

23-25 March 1998

Scott M. Brown, Eugene Santos Jr., and Sheila B. Banks

Department of Electrical & Computer Engineering
Air Force Institute of Technology
2950 P Street, Wright-Patterson AFB, OH 45433-7765

Supporting Incremental Knowledge Elicitation in Decision-Theoretic Systems

Scott M. Brown
Depart. of Electrical and Computer Engineering
Air Force Institute of Technology
Wright-Patterson AFB, OH 45433-7765
sbrown@afit.af.mil

Eugene Santos Jr.
Computer Science and Engineering
University of Connecticut
Storrs, CT 06269-3155
eugene@eng2.uconn.edu

Sheila B. Banks
Depart. of Electrical and Computer Engineering
Air Force Institute of Technology
Wright-Patterson AFB, OH 45433-7765
sbanks@afit.af.mil

Introduction

Knowledge elicitation continues to be a bottleneck to constructing decision-theoretic systems. Most knowledge representations for these systems require complete knowledge of the domain before the systems become useable. Methodologies and techniques for incremental elicitation of knowledge in support of users' current goals is desirable. A primary goal of our research is to develop a comprehensive software engineering, knowledge engineering, and knowledge elicitation methodology for Symbiotic Information Reasoning and Decision Support (Banks *et al.* 1997). To that end, in this position paper we briefly discuss Bayesian knowledge bases, a probabilistic knowledge representation allowing for incomplete specification of knowledge. We describe how Bayesian knowledge bases along with an intelligent interface agent are used in an expert system shell called PESKI to support incremental knowledge elicitation.

Bayesian Knowledge Bases

To support the design of decision-theoretic systems, we desire to have a knowledge representation that supports modeling uncertainty and is flexible, intuitive, and mathematically sound. A Bayesian knowledge base (abbrev. BKB) is a probabilistic knowledge representation meeting the preceding qualities (Santos Jr. & Santos 1996). A BKB supports theoretically sound and consistent probabilistic inference — even with incomplete knowledge — with the intuitiveness of "if-then" rule specification. The representation is similar to Bayesian Networks (Pearl 1988); it is a directed graph capable of representing uncertainty in knowledge via probabilistic relationships between random variables. However, Bayesian networks do not allow for incompleteness.

The PESKI Environment

PESKI (Probabilities, Expert Systems, Knowledge, and Inference) is an integrated probabilistic knowledge-based expert system shell utilizing Bayesian knowledge bases as its knowledge representation. PESKI provides users with knowledge acquisition (Santos Jr., Banks, & Banks 1997), verification and validation (Santos Jr., Gleason, & Banks 1997; Bawcom 1997), data mining, and inference engine tools (Shimony, Domshlak, & Santos Jr. 1997), each capable of operating in various communication modes. For more information on PESKI, see the United States Air Force Institute of Technology's Artificial Intelligence Laboratory web site[1].

PESKI supports incremental knowledge elicitation in a number of ways (Santos Jr., Banks, & Banks 1997). During knowledge acquisition, the user is alerted to any inconsistencies in the BKB knowledge representation. For example, if the user attempts to add a rule that creates a cycle in the knowledge base, PESKI will display an error message to the user. We use a test case-based approach to knowledge base verification and validation (Santos Jr., Gleason, & Banks 1997). The user submits test cases by providing evidence and a proposed answer(s) and PESKI determines via inference if, with the given knowledge, the answer specified

[1] `http://www.afit.af.mil/Schools/EN/ENG/LABS/AI/`

in the test case could be obtained given the evidence. If PESKI can not, incompleteness exists and additional knowledge must be elicited. PESKI supports the correction of any incompleteness found using a graphical incompleteness tool (Bawcom 1997). Figure 1 shows an example of the use of this tool in PESKI. The tool uses data visualization of the BKB and data mining to assist the user in eliciting the needed knowledge.

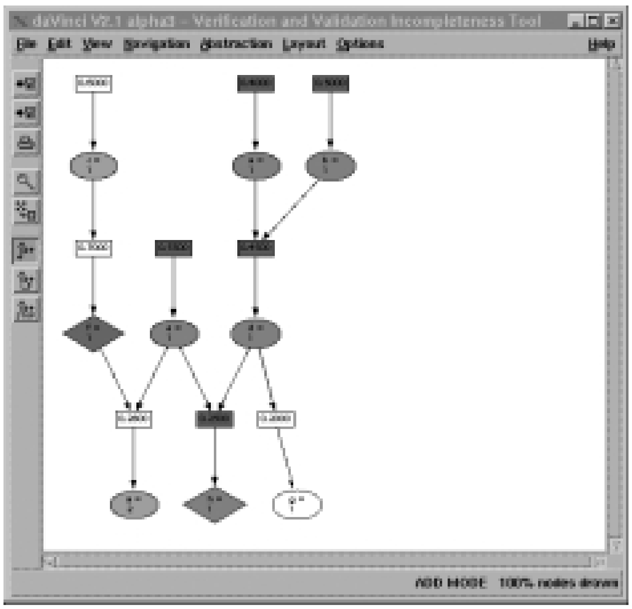

Figure 1: PESKI's Graphical Incompleteness Tool.

We have integrated an intelligent interface agent into PESKI (Harrington, Banks, & Santos Jr. 1996a; 1996b; Harrington & Brown 1997). The overall goal of the agent is to offer timely, beneficial assistance to the user as he/she interacts with PESKI. To accomplish this goal, an accurate cognitive model of the user is maintained (Brown *et al.* 1998). The agent is currently capable of offering assistance for tool, communication mode, and knowledge base use. We are currently researching expanding the agent's user model to allow the agent to elicit information from the user based on what goals he/she is trying to achieve, his/her preferences, and past actions. To that end, we are adding domain knowledge of BKBs to the interface agent's user model.

References

Banks, S. B.; Stytz, M. R.; Santos Jr., E.; and Brown, S. M. 1997. User modeling for military training: Intelligent interface agents. In *Proceedings of the 19th Interservice/Industry Training Systems and Education Conference.*

Bawcom, D. 1997. An incompleteness handling methodology for validation of bayesian knowledge bases. Master's thesis, Air Force Institute of Technology.

Brown, S. M.; Santos Jr., E.; Banks, S. B.; and Oxley, M. 1998. Using explicit requirements and metrics for interface agent user model correction. In *Proceedings of the Second International Conference on Autonomous Agents (Agents '98).* to appear.

Harrington, R. A., and Brown, S. M. 1997. Intelligent interface learning with uncertainty. In Santos Jr., E., ed., *Proceedings of the Eighth Midwest Artificial Intelligence and Cognitive Science Conference*, 27–34. AAAI Press.

Harrington, R. A.; Banks, S.; and Santos Jr., E. 1996a. Development of an intelligent user interface for a generic expert system. In Gasser, M., ed., *Online Proceedings of the Seventh Midwest Artificial Intelligence and Cognitive Science Conference.* Available at http://www.cs.indiana.edu/event/maics96/.

Harrington, R. A.; Banks, S.; and Santos Jr., E. 1996b. GESIA: Uncertainty-based reasoning for a generic expert system intelligent user interface. In *Proceedings of the 8th IEEE International Conference on Tools with Artificial Intelligence*, 52–55.

Pearl, J. 1988. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference.* San Mateo, CA: Morgan Kaufmann.

Santos Jr., E., and Santos, E. S. 1996. Bayesian knowledge-bases. Technical Report AFIT/EN/TR96-05, Department of Electrical and Computer Engineering, Air Force Institute of Technology, Wright-Patterson AFB, OH.

Santos Jr., E.; Banks, D. O.; and Banks, S. B. 1997. Mack: A tool for acquiring consistent knowledge under uncertainty. In *Proceedings of the AAAI Workshop on Verification and Validation of Knowledge-Based Systems*, 23–32.

Santos Jr., E.; Gleason, H. T.; and Banks, S. B. 1997. Bval: Probabilistic knowledge-base validation. In *Proceedings of the AAAI Workshop on Verification and Validation of Knowledge-Based Systems*, 13–22.

Shimony, S. E.; Domshlak, C.; and Santos Jr., E. 1997. Cost-sharing heuristic for bayesian knowledge-bases. In *Proceedings of the Thirteenth Conference on Uncertainty in Artificial Intelligence*, 421–428.